

Claims

What is claimed is:

- 5 1. A database management system comprising:
 a mainline database system that makes modifications to data in the database management system using a write-ahead logging protocol; stores data on a first set of storage volumes and stores log records on a second set of storage volumes; restores consistency between the log records and the data during a restart, and while a backup
10 system lock is held by a backup utility, continues updating objects except for suspending actions that change an external file system catalog, suspending writing updates of objects that extend across a storage volume boundary; and freezing a REDO log point in checkpoint information while the backup system lock is taken by the backup utility.
- 15 2. The database management system of claim 1 further comprising a backup utility that obtains the backup system lock before starting a backup process; copies the first set of storage volumes to a first set of backup volumes; records information identifying the first set of backup volumes in a dataset.
- 20 3. The database management system of claim 2, wherein the backup utility copies the second set of storage volumes to the backup medium; and records backup volume information for the second set of storage volumes in the dataset.
- 25 4. The database management system of claim 2 further comprising a restore utility that performs a point-in-time recovery using the data from the first set of backup volumes, a user specified point-in-time, and the logs on the second set of storage volumes.

5. The database management system of claim 4 wherein the restore utility marks a first object as recovery-pending when a log record identifies the first object as having been updated without log records so that subsequent restoration of the first object can be made from an image copy.

5

6. The database management system of claim 1 wherein the mainline database system writes log records to identify objects that have been updated without log records and writes log records to identify objects that have been created, extended and/or deleted.

10

7. The database management system of claim 1 wherein the mainline database system stores checkpoint information periodically and the backup utility records a log apply starting point corresponding to a last checkpoint information storage point.

15 8. The database management system of claim 1 wherein the mainline database system obtains the backup system lock before updating objects that change an external file system catalog or that extend across a storage volume boundary.

9. A database management system, that performs a backup without suspending updates and the backup can be restored using a write-ahead logging protocol restart, comprising:
means for modifying data in the database management system using a write-ahead logging protocol;

5 means for restoring consistency between the log records and the data during a restart;

means for storing data on a first set of storage volumes and storing log records on a second set of storage volumes;

10 means for freezing a REDO log point in checkpoint information while a backup system lock is taken; and

means for continuing to update the data while the backup system lock is taken, except for suspending actions that change an external file system catalog, and except for suspending writing updates of objects that extend across a storage volume boundary.

15 10. The database management system of claim 9 further comprising:

means for obtaining the backup system lock before starting a backup process; for copying the first set of storage volumes to a first set of backup volumes; and for recording information identifying the first set of backup volumes in a recovery control dataset and in an external file system's control dataset.

20

11. The database management system of claim 10 further comprising means for copying the second set of storage volumes to a second set of backup volumes; and means for recording information identifying the second set of backup volumes in the recovery control dataset and in the external file system's control dataset.

25

12. The database management system of claim 10 further comprising means for restoring data from the first set of backup volumes to the first set of storage volumes and for performing a point-in-time recovery using a user specified point-in-time, and the logs on the second set of storage volumes.

5

13. The database management system of claim 12 wherein the means for restoring data further comprises means for marking a first object as recovery-pending when a log record identifies the first object as having been updated without log records so that subsequent restoration of the first object can be made from an image copy; processing log records identifying a second object which has been newly created by allocating space for the second object; processing log records identifying a third object which has been newly extended by allocating additional space for the third object; processing log records identifying a fourth object which has been deleted by freeing space for the fourth object; and means for setting a mode to indicate that the point-in-time recovery has completed.

15

14. The database management system of claim 9 further comprising means for writing log records to identify objects that have been updated without log records and writing log records to identify objects that have been created, extended or deleted.

20

15. The database management system of claim 10 further comprising means for storing checkpoint information periodically and the means for backing up the data further comprises means for recording a log apply starting point corresponding to a last checkpoint information storage point.

25

16. The database management system of claim 9 further comprising means for obtaining the backup system lock before updating objects that change an external file system catalog or that extend across a storage volume boundary.

17. A method of operating a database management system comprising the steps of:
modifying data in the database management system using a write-ahead logging
protocol;
restoring consistency between log records and the data during a restart;
5 storing data on a first set of storage volumes and storing log records on a second
set of storage volumes;
freezing a REDO log point in checkpoint information while a backup system lock
is taken; and
continuing to update the data while the backup system lock is taken, except for
10 suspending actions that change an external file system catalog, and except for suspending
writing updates of objects that extend across a storage volume boundary.
18. The method of claim 17 further comprising the step of backing up the data and
wherein the step of backing up the data further comprises obtaining the backup system
15 lock and after obtaining the backup system lock, copying the first set of storage volumes
to a first set of backup volumes.
19. The method of claim 18 further wherein the step of backing up the data further
comprises the step of recording information identifying the first set of backup volumes in
20 a recovery control dataset and in an external file system's control dataset.

20. The method of operating the database management system of claim 18 further comprising backing up log records, after backing up the data, by copying the second set of storage volumes to a second set of backup volumes.

5 21. The method of operating the database management system of claim 18 further comprising the steps of restoring data from the first set of backup volumes to the first set of storage volumes and performing a point-in-time recovery using a user specified point-in-time, and the logs on the second set of storage volumes.

10 22. The method of operating the database management system of claim 21 wherein the step of performing a point-in-time recovery further comprises the steps of marking a first object as recovery-pending when a log record identifies the first object as having been updated without log records so that subsequent restoration of the first object can be made from an image copy; processing log records identifying a second object which has been
15 newly created by allocating space for the second object; processing log records identifying a third object which has been newly extended by allocating additional space for the third object; processing log records identifying a fourth object which has been deleted by freeing space for the fourth object; and setting a mode to indicate that the point-in-time recovery has completed.

20

23. The method of operating database management system of claim 17 further comprising writing log records to identify objects that have been updated without log records and writing log records to identify objects that have been created, extended or
25 deleted.

25

24. The method of operating the database management system of claim 18 further comprising storing checkpoint information periodically and the step of backing up the data further comprises recording a log apply starting point corresponding to a last checkpoint information storage point.

5

25. The method of operating the database management system of claim 17 wherein the step of continuing to update the data while a backup system lock is taken further comprises obtaining the backup system lock before updating objects that change an external file system catalog or that extend across a storage volume boundary.

10

26. An article of manufacture comprising computer usable media including at least one computer program recorded therein that is capable of causing a computer system to perform a method of operating a database management system comprising the steps of:

15 modifying data in the database management system using a write-ahead logging protocol;

 restoring consistency between log records and the data during a restart;

 storing data on a first set of storage volumes and storing log records on a second set of storage volumes;

20 freezing a REDO log point in checkpoint information while a backup system lock is taken; and

 continuing to update the data while the backup system lock is taken, except for suspending actions that change an external file system catalog, and except for suspending writing updates of objects that extend across a storage volume boundary.

25 27. The article of manufacture of claim 26 wherein the method further comprises the step of backing up the data and wherein the step of backing up the data further comprises obtaining the backup system lock and after obtaining the backup system lock, copying the first set of storage volumes to a first set of backup volumes.

28. The article of manufacture of claim 27 wherein the step of backing up the data further comprises the step of recording information identifying the first set of backup volumes in a control dataset.

5

29. The article of manufacture of claim 34 wherein the method further comprises backing up log records, after backing up the data, by copying the second set of storage volumes to a second set of backup volumes.

10 30. The article of manufacture of claim 29 wherein the method further comprises recording information identifying the second set of backup volumes in the recovery control dataset and in the external file system's control dataset.

15 31. The article of manufacture of claim 27 wherein the method further comprises the steps of restoring data from the first set of backup volumes to the first set of storage volumes and performing a point-in-time recovery using a user specified point-in-time, and the logs on the second set of storage volumes.

20 32. The article of manufacture of claim 31 wherein the step of performing a point-in-time recovery further comprises the steps of marking a first object as recovery-pending when a log record identifies the first object as having been updated without log records so that subsequent restoration of the first object can be made from an image copy; processing log records identifying a second object which has been newly created by allocating space for the second object; processing log records identifying a third object
25 which has been newly extended by allocating additional space for the third object; processing log records identifying a fourth object which has been deleted by freeing space for the fourth object; and setting a mode to indicate that the point-in-time recovery has completed.

33. The article of manufacture of claim 26 wherein the step of continuing to update the data while a backup system lock is taken further comprises obtaining the backup system lock before updating objects that change an external file system catalog or that extend
5 cross a storage volume boundary.